

High Precision Encoding and Decoding of Video Images

TECHNICAL FIELD

This invention relates to video compression, and more particularly to improved quality video compression based on novel improvements to MPEG-like encoding and decoding systems.

BACKGROUND

MPEG Background

MPEG-2 and MPEG-4 are international video compression standards defining a video syntax that provides an efficient way to represent image sequences in the form of more compact coded data. The language of the coded bits is the "syntax." For example, a few tokens can represent an entire block of samples (*e.g.*, 64 samples for MPEG-2). Both MPEG standards also describe a decoding (reconstruction) process where the coded bits are mapped from the compact representation into an approximation of the original format of the image sequence. For example, a flag in the coded bitstream signals whether the following bits are to be preceded with a prediction algorithm prior to being decoded with a discrete cosine transform (DCT) algorithm. The algorithms comprising the decoding process are regulated by the semantics defined by these MPEG standards. This syntax can be applied to exploit common video characteristics such as spatial redundancy, temporal redundancy, uniform motion, spatial masking, *etc.* In effect, these MPEG standards define a programming language as well as a data format. An MPEG decoder must be able to parse and decode an incoming data stream, but so long as the data stream complies with the corresponding MPEG syntax, a wide variety of possible data structures and compression techniques can be used. It is also possible to carry the needed semantics within an alternative syntax.

These MPEG standards use a variety of compression methods, including intraframe and interframe methods. In most video scenes, the background remains relatively stable while action takes place in the foreground. The background may move, but a great deal of the scene is redundant. These MPEG standards start compression by creating a reference frame called an "Intra" frame or "I frame". I frames are compressed without reference to other frames and thus contain an entire frame of video information. I frames provide entry points into a data bitstream for random access, but can only be moderately compressed. Typically, the data representing

I frames is placed in the bitstream every 12 to 15 frames. Thereafter, since only a small portion of the frames that fall between the reference I frames are different from the bracketing I frames, only the image differences are captured, compressed, and stored. Two types of frames are used for such differences – Predicted or P frames, and Bi-directional Interpolated or B frames.

5 P frames generally are encoded with reference to a past frame (either an I frame or a previous P frame), and, in general, are used as a reference for subsequent P frames. P frames receive a fairly high amount of compression. B frames provide the highest amount of compression but require both a past and a future reference frame in order to be encoded. Bi-directional frames are never used for reference frames.

10 Macroblocks are regions of image pixels. For MPEG-2, a macroblock is a 16x16 pixel grouping of four 8x8 DCT blocks, together with one motion vector for P frames, and one or two motion vectors for B frames. Macroblocks within P frames may be individually encoded using either intra-frame or inter-frame (predicted) coding. Macroblocks within B frames may be individually encoded using intra-frame coding, forward predicted coding, backward predicted
15 coding, or both forward and backward (*i.e.*, bi-directionally interpolated) predicted coding.

After coding, an MPEG data bitstream comprises a sequence of I, P, and B frames. A sequence may consist of almost any pattern of I, P, and B frames (there are a few minor semantic restrictions on their placement). However, it is common in industrial practice to have a fixed pattern (*e.g.*, IBBPBBPBBPBBPBB).

20 It has been known for some time that computation is reduced when determining motion vectors by utilizing a hierarchical motion search. For example, the MPEG algorithms attempt to find a match between “macroblock” regions. MPEG-type and other motion compensated DCT (discrete cosine transform) coders attempt to match each macroblock region in a current frame with a position in a previous frame (P frame) or previous and
25 subsequent frame (B frame). However, it is not always necessary to find a good match, since MPEG can code a new macroblock as a fresh stand-alone (“intra”) macroblock in this case without using previous or subsequent frames. In such motion compensated DCT systems, one macroblock motion vector is needed for each macroblock region for MPEG-2. In MPEG-4, a set of 4 motion vectors, corresponding to one vector for each 8x8 region (*i.e.*, 4 vectors per
30 macroblock) is also an optional coding mode.

MPEG Precision

The reference MPEG-2 and MPEG-4 video codec implementations utilize the following encoding methodology:

a) When converting from RGB to YUV color space, only the number of bits that will be coded are kept (for example, MPEG-2 is limited to 8 bits in coding, and thus the YUV values are also limited to 8 bits).

b) When encoding and decoding, only the number of bits that have been coded are preserved, with careful rounding being applied to reduce artifacts.

c) When converting back to RGB, the precision is limited due to the limitations of the number of bits which were preserved (such as 8 bits maximum for MPEG-2).

FIG. 1 is a block diagram of a prior art MPEG-2 reference video encoding method. RGB input frames 102 coded in 8 bits/pixel per color are applied to an RGB-to-YUV converter 104, which is purposely limited to 8 bits of precision per color on its output. The result is applied to a DCT function 106, then to a quantizer function 108, then to an inverse DCT function 110, with the final output 212 being stored at the same precision as the input data.

MPEG-4's reference video coder is implemented with the same method, although the intermediate precision can be extended up to 12 bits (although the VLC – variable length coding – tables do not support use of the full range).

Techniques for randomly dithering the limited precision values (8 bits per color component maximum in MPEG-2) are utilized to reduce the apparent visibility of step changes. However, noise and artifacts in coding are created due to this dither, and are also created due to the use of limited intermediate processing precision.

In addition to limited intermediate processing precision, MPEG-2 and MPEG-4 allow the inverse DCT (IDCT) algorithm used during encoding (often implemented in high precision floating point representation) to differ slightly from the IDCT algorithm used during decoding. This is known as “IDCT mismatch”. IDCT mismatch causes an unpredictable gradual drift in the signal away from the intended decoding values. This is conventionally reduced by use of random dither of the low order bit in the IDCT highest frequency (7th harmonic for the typical 8x8 DCT block size used in MPEG-2 and MPEG-4). Such dithering adds additional noise and artifacts to the signal.

FIG. 2 is a block diagram of a prior art MPEG-2 reference video decoding method. An encoded input bitstream 202 is applied to a dequantizer function 204 having a limited precision that matches the precision of the input bitstream (typically 8 bits for MPEG-2). The result is applied to an IDCT function 206 (which may not match the IDCT function 110 of the encoder), which output signed 8-bit values 208. This output comprises either an I frame 210, or is combined either with data from a previous frame 212 or a subsequent frame 214 (both at the same precision) to generate a new frame 216. Thus, the MPEG-2 decoding process limits intermediate processing precision to a maximum of 8 bits. Similarly, the intermediate processing precision for MPEG-4 video decoding is also limited to the number of bits used in encoding (a maximum of 12 bits, but often set to be 8 bits).

Limited precision in MPEG-2 and MPEG-4 also limits dynamic range (*i.e.*, the number of levels of lighting that can be represented for an image) and contrast range (*i.e.*, the number of distinct levels assigned to image regions of similar contrast). Accordingly, the encoding and decoding methods used in MPEG-2 and MPEG-4 reduce the potential quality of output, decompressed images compared to the original input images. The present invention addresses these limitations.

SUMMARY

The invention is directed to improved quality video compression based on novel improvements to MPEG-like encoding and decoding systems. In one aspect, the invention provides a method for significantly improving image quality from MPEG-style video coding by preserving a higher number of bits during intermediate encoding and decoding processing steps. Surprisingly, this improvement in quality does not result in a proportionally greater overall number of bits required to encode a sequence of images. Further, the problems of IDCT mismatch can be eliminated by exactly matching the IDCT function numerical algorithm of the decoder to the IDCT function numerical algorithm used for the decoding portion of the encoder. Eliminating the IDCT mismatch allows an increase in compression ratios by reducing the number of required I frames.

In another aspect, the invention includes application of high precision compression to wide dynamic range images by extending the range of the "quantization parameter" or "QP". Dynamic range extension uses a low QP for dark regions and a high QP for bright regions. The extension of QP may be accomplished either by increasing the range of QP directly, or indirectly through a non-linear transformation (such as a function or lookup table) which maps a small range of QP values to a wide range of output QP values for dividing (compression) and multiplying (decompression).

In another aspect, the invention includes application of extended intermediate processing precision and an extended QP range to reduced contrast regions of an image to extend the precision with which the low-contrast portions are compression coded. A low QP can be utilized with low-contrast (often distant) regions, whereas a high QP is utilized for high contrast (usually foreground) regions. In this way, for example, if a camera sensor has sufficient bit resolution, distant objects in the sky or on the ground on a hazy day can be fully distinguished when decompressed.

In particular, in one aspect the invention includes a method for compressing a sequence of digitized video images including a sequence of frames represented at a first precision in a first color space, the method including transforming the sequence of frames to a representation in a second color space at a second precision greater than the first precision, and performing subsequent encoding steps at the second precision to create a compressed output.

Another aspect of the invention is a method for compressing a sequence of digitized video images including a sequence of frames represented at a first precision in a first color space, the method including transforming the sequence of frames to a representation in a second color space at a second precision greater than the first precision; performing a motion compensated discrete cosine transform at the second precision on the sequence of frames to produce a first encoded output; quantizing the first encoded output at the second precision to create a quantized output; performing an inverse discrete cosine transform at the second precision on the quantized output to produce a compressed output at the second precision.

In another aspect, the invention includes a method for compressing and decompressing a sequence of digitized video images including a sequence of frames represented at a first precision in a first color space, the method including encoding the sequence of frames to create a compressed bitstream, and performing decoding steps on the compressed bitstream at a second precision greater than the first precision.

Another aspect of the invention is a method for compressing and decompressing a sequence of digitized video images including a sequence of frames represented at a first precision in a first color space, the method including transforming the sequence of frames to a representation in a second color space at a second precision greater than the first precision; performing subsequent encoding steps at the second precision to create a compressed bitstream; and decoding the compressed bitstream by dequantizing the compressed bitstream at the second precision to create a dequantized output, applying an inverse discrete cosine transform at the second precision on the dequantized output to produce a decompressed output, and generating image frames at the second precision from the decompressed output.

Yet another aspect of the invention is a method for compressing and decompressing a sequence of digitized video images including a sequence of frames represented at a first precision in a first color space, the method including transforming the sequence of frames to a representation in a second color space at a second precision greater than the first precision; performing subsequent encoding steps at the second precision to create a compressed bitstream, including applying an inverse discrete cosine transform numerical algorithm; and performing subsequent decoding steps at the second precision on the compressed bitstream, including applying a matching inverse discrete cosine transform numerical algorithm.

Another aspect of the invention is a method for increasing the average compression ratio during compression of a sequence of digitized video images including a sequence of frames to a compressed bitstream that includes I frames, the method including encoding the sequence of frames by applying an inverse discrete cosine transform numerical algorithm that
5 matches the inverse discrete cosine transform numerical algorithm applied during a subsequent decoding process, thereby generating a compressed bitstream having increased spacing between I frames.

Another aspect of the invention is a method for compressing a sequence of digitized video images having at least one of a wide dynamic range or wide contrast range, the video
10 images including frames represented at a first precision, each frame including a plurality of macroblocks defining regions, the method including determining a quantization parameter, QP, having codes that represent either a direct correspondence to possible values or an extended range of possible values; determining a QP value for each macroblock of each frame; compressing each frame at a second precision higher than the first precision to create
15 a compressed frame, such compressing including applying the determined QP values for such frame to reduce the number of bits required to encode such frame; associating QP codes with the determined QP values used during compressing; and outputting each compressed frame and the associated QP codes.

The invention includes corresponding computer program implementations and
20 apparatus implementations.

The details of one or more embodiments of the invention are set forth in the accompanying drawings and the description below. Other features, objects, and advantages of the invention will be apparent from the description and drawings, and from the claims.

DESCRIPTION OF DRAWINGS

FIG. 1 is a block diagram of a prior art MPEG-2 reference video encoding method.

FIG. 2 is a block diagram of a prior art MPEG-2 reference video decoding method.

5 FIG. 3 is a block diagram of an MPEG-like encoding method in accordance with the present invention.

FIG. 4 is a block diagram of an MPEG-like decoding method in accordance with the present invention.

FIG. 5 is a flowchart that summarizes a preferred method of extending dynamic range and/or contrast during image compression.

10 Like reference symbols in the various drawings indicate like elements.

DETAILED DESCRIPTION

Higher Precision Intermediate Image Processing

The limited precision assumptions and techniques within MPEG-2 and MPEG-4 are designed to minimize the amount of memory needed for storing I, B, and P frames. However, frame memory is presently quite affordable. The present invention is based in part on the discovery that a significantly improved image quality can be achieved from MPEG-style video coding by preserving a higher number of bits during intermediate encoding and decoding processing steps. Surprisingly, this improvement in quality does not result in a proportionally greater overall number of bits required to encode a sequence of images. In fact, the number of bits is usually reduced using the present invention.

FIG. 3 is a block diagram of an MPEG-like encoding method in accordance with the present invention. RGB input frames 302 are applied to an RGB-to-YUV converter 304. Because modern frame generation devices (*e.g.*, video cameras and high-definition film scanners) can output higher color range (*e.g.*, 10 bits/pixel per color) images, such devices are preferred as the input source. However, the input source may be a conventional MPEG-2 or MPEG-4 8-bit device. The result is applied to a DCT function 306, which preserves more bits of precision (*e.g.*, 16 bits) than are present in the original input signal. The output of the DCT function 306 is applied to a quantizer function 308, then to an IDCT function 310, which again preserves more bits of precision (*e.g.*, 16 bits) than are present in the original input signal (such as 16-bits, signed for P and B frames, unsigned for I frames, 16 bits being a convenient representation for digital hardware and software systems). The final output 312 is typically a YUV signal stored at the same precision as the intermediate processing precision. An important characteristic of such increased precision in the output is that it permits improved prediction of subsequent P and B frames.

The concept of IDCT mismatch which is embodied in MPEG-2 and MPEG-4 video coding is based on the assumption that the computation used for decoding may differ from the computation used for the decoding portion of encoding. As mentioned above, this mismatch will cause drift, even in the presence of the high-harmonic low-bit DCT mismatch dither (which also adds noise). It is common practice in MPEG-2 to place I frames near each other (approximately every half second) to re-correct this drift, and to limit the extent of the error. However, I frames are relatively inefficient, usually costing about 3 times as many bits

as P frames, and 5 times as many bits as B frames. I frames also form points of restart and reference during “tune-in” to a motion picture sequence. However, the frequency of their occurrence on stored media could usefully be increased to several seconds (*e.g.*, in the range of about 1 to 5 seconds, on average) to improve efficiency, were it not for the IDCT mismatch.

The IDCT mismatch can be eliminated by exactly matching the numerical algorithms of the IDCT function of the decoder to those of the IDCT function used for the decoding portion of the encoder. Any precision limitations in these matching IDCT functions are automatically corrected at each P frame due to the natural feedback mechanism that arises in going from one P frame to the next via coding its difference signal (which includes the difference of IDCT precision limitations). “Exact matching” of the numerical algorithms means that those portions of the algorithms that transform input to output should apply the same definitions for multiply and add functions, the same numeric representations, the same precision, *etc.* However, the numerical algorithms need not be identical as to computer program or integrated circuit implementation. Thus, for example, different computer languages and binary generating modes (*e.g.*, interpreted vs. compiled) may be used.

Thus, high quality coding can be achieved by sufficient precision in the IDCT function. However, the IDCT function need not require very high precision. For example, in the MPEG-4 video reference software, a double-width floating point (64-bit) IDCT implementation is used. This is completely unnecessary, since a 16-bit IDCT implementation is sufficient to provide the improvements necessary for coding up to 12-bit dynamic range. Encoder and decoder IDCT implementations (exactly matching) greater than 16 bits, such as 32-bit floating point implementations, can be used to extend the dynamic range up to 16 bits (which results in greater than a 1 billion-to-one dynamic range in less than tenth-percent logarithmic steps, exceeding the limits of human vision). Thus, by exactly matching the encoder and decoder IDCT implementations, the present invention greatly reduces the amount of computation needed for the IDCT implementations while eliminating the problems of IDCT mismatch. Further, and counter-intuitively, using exactly matching encoder and decoder IDCT implementations actually increases overall efficiency (*i.e.*, higher average compression ratio) even with increased intermediate precision, since bit-costly I frames can be spaced further apart in time (*e.g.*, in the range of about 1 to 5 seconds, on

average). Indeed, I frames can be spaced apart by virtually unlimited times, limited only by the desire to be able to jump into the middle of a program or to correct errors generated from a lossy distribution channel.

FIG. 4 is a block diagram of an MPEG-like decoding method in accordance with the present invention. A high-precision encoded input bitstream 402 is applied to a dequantizer function 404 having an "intermediate" processing precision that matches the precision of the input bitstream. The result preferably is applied to an IDCT function 406 that is an exact match to the IDCT function 310 of the corresponding encoder. The IDCT function 406 outputs signed values 408 of the same intermediate precision as all prior internal processing steps (*e.g.*, 16 bits). This output comprises either an I frame 410, or is combined either with data from a previous frame 412 or a subsequent frame 414 (both at the same precision) to generate a new frame 416.

In addition, all forms of dither should be eliminated, thus reducing noise and artifacts. In particular, dither from rounding (except at full precision, such as 16 bits – *i.e.*, round the 17th bit), and dither of the low bit of the high harmonic from IDCT mismatch, should both be eliminated. Also, in the preferred embodiment, the additional intermediate pixel precision is used during any final color space conversion step (*e.g.*, YUV to RGB or other conversions, such as YUV 4:2:0 to YUV 4:2:2, for viewing, utilizing, or storing the converted image) during decoding, rounding only at the final step.

It should be noted that the high precision techniques shown in FIG. 3 and FIG. 4 may be used to encode and subsequently decode a standard precision input (*e.g.*, the 8-bit precision input used by MPEG-2). While the result is not as high in quality as encoding and decoding a higher precision input, the result will still be an improvement over present MPEG encoding and decoding. Further, both the encoding and decoding process can be improved by increased precision during intermediate processing and storage. Such precision need not be identical to gain improvement, but the improvement is optimized when the decoding portions of encoding and decoding exactly match in precision and numerical algorithm.

Following is a summary of the preferred embodiment of the present method for improving compressed image quality:

1) Preserve more bits of precision during intermediate processing than the precision of the input (*e.g.*, preserve more bits of precision from the RGB to YUV conversion step during encoding, and preserve more bits of precision from the IDCT step).

2) Store the increased intermediate precision result.

5 3) Optionally, utilize an exactly matching IDCT implementation in the encoder and decoder.

4) Optionally, eliminate all forms of dither.

5) Utilize the additional pixel precision during the final color space conversion step during decoding, rounding only at the final step.

10 *Dynamic Range and Contrast Extension*

The present inventor has previously discussed the concept of gradually increasing the colorimetric and dynamic range of pixel-based image representations. See, *e.g.*, “The Use of Logarithmic and Density Units for Pixels” by Gary Demos, presented at the October 1990 SMPTE conference, and published in in the SMPTE Journal (Oct. 1990, vol. 100, no. 10).

15 See also “An Example Representation For Image Color And Dynamic Range Which Is Scalable, Interoperable, and Extensible” by Gary Demos, presented at the October 1993 SMPTE conference and published in the proceedings and preprints.

The use of a logarithmic representation for pixel values has many benefits. For example, the YUV coding methodology computes U as R-Y, and V as B-Y. In a logarithmic representation, U becomes R/Y and V becomes B/Y, both of which are “isoluminant” in the terminology coined in the 1993 SMPTE paper, “An Example Representation For Image Color And Dynamic Range Which Is Scalable, Interoperable, and Extensible”, referenced above. That is, both U and V channels contain no signal under variation of illumination if they have a constant hue. This provides for high efficiency in coding color using the U and V channels. Further, this hue coding efficiency is obtained over a very wide dynamic range of overall brightness. A logarithmic representation also provides for easy methods of system measurement and calibration, as well as being perceptually uniform over a very wide range of brightness.

Table 1 indicates the range and tonal definition capability of various numbers of bits in the range of 9 to 14 bits/pixel. It can be seen from this table that the range of human

vision, which spans a range of approximately 10,000,000-to-1 in brightness range, color range, and tonal definition (in 1/8 percent steps), can be approached using a precision of less than 16 bits using a logarithmic representation.

Total Contrast Range	Number of Required Bits			
	1% Steps	.5% Steps	.25% Steps	.125% Steps
165 : 1	9	10	11	12
30000 : 1	10	11	12	13
700,000,000 : 1	11	12	13	14

Number of Bits Required Using Whole, Half, Quarter, & Eighth Percent Logarithmic Steps
Table 1

Current high quality electronic cameras and projectors are capable of operating at approximately 10-bits/pixel per color of dynamic and color range. For example, the Polaroid/Philips LDK9000 camera provides a low-noise image with a wide color range. This camera has an 11 micron CCD pixel size and a full well electron count of approximately 25,000 electrons. Larger sensor sizes are very feasible, with the potential to increase the full well electron count to hundreds of thousands or millions of electrons. On the image projection side, micro-mirror projectors with 10-bit gamma input are able to achieve a 1000:1 dynamic range with reasonable tonal distinction, thereby approaching the quality of movie film. Although the best films can achieve wider dynamic range (approximately 3000:1) with high tonal fidelity and broad color range, it is now possible to foresee digital imaging and presentation rivaling and eventually exceeding this performance.

As noted above, the quality of moving image compression can be significantly improved by retaining extended precision during intermediate processing. The same mechanism can also be utilized to greatly extend the dynamic range of image information that can be efficiently compressed. For example, if 14 bits of intermediate precision are retained, then this bit range can represent a dynamic range of 700,000,000:1 at 1/8% logarithmic steps.

It is also useful to make the additional observation, not covered in the referenced papers, that brightness distinction is local. Thus, it is not possible to distinguish small brightness variations in a dark shadow area that is immediately adjacent to a very bright object. Thus, it is only necessary to retain tonal and dynamic range distinction with respect to the local brightness in that same region of an image. A different portion of the image,

however, might have dark shadows, and might be far enough from the bright region that substantial distinction of detail is seen, requiring corresponding detail in the tonal range of the local brightness representation.

These concepts, combined with the use of extended precision during intermediate processing, can be applied to moving image compression. In particular, once the intermediate precision which is maintained within the compression system is extended to higher numbers of bits, such as 13, 14, 15, or 16, then this extended precision is also capable of representing wide dynamic range images. Further, in order to obtain efficient compression, the perceptual limitation of tonal distinction in bright areas versus the expanded distinction in shadows can be utilized.

In MPEG-like compression systems, the tonal distinction is determined by the "quantization parameter" or "QP". QP is divided into the DCT frequency coefficients to reduce the number of bits required to code a moving image stream. During decoding, the QP is multiplied times the DCT coefficients prior to computing the IDCT. Thus, although QP is applied in frequency (DCT) space, it still represents a tonal precision parameter.

In light of the regional distinction characteristics described above, a high QP can be utilized in areas of high brightness without visible loss of clarity or tone, since tonal distinction is relative to full brightness. However, in the dark image shadow regions, a low QP must be utilized to provide for fine tonal precision.

It is common in MPEG-1, MPEG-2, and MPEG-4 to utilize a linear QP factor in the range of 1 to 32. While this is suitable for an 8-bit dynamic range, such as is provided by MPEG-2, this range is insufficient for higher numbers of bits (such as 10-bits or 12-bits), or for wider dynamic range. In MPEG 2 and MPEG-4, it is possible to vary QP from one macroblock to the next. This is normally the mechanism by which bit rate is adjusted to maintain a constant bit rate. A higher QP produces fewer coded bits, while a lower QP produces more coded bits. Thus, varying QP in the range of 1 to 32 is all that is required in order to maintain a given constant bit rate in a limited precision system such as the 8-bit capability of MPEG-1 or MPEG-2. However, for 10-bit or 12-bit precision, as in MPEG-4, if the amount of scene change is widely varying (high scene stress), and a low constant bit rate is required, then a QP range of 32 possible values may be insufficient. A QP range of 32 values for 10 bits is equivalent to a QP range of 8 values for 8 bits, being only a quarter of

the range available to 8-bit coding systems such as MPEG-2. For 12-bit encoding systems, such as MPEG-4, a range of 32 values is equivalent to a QP range of 2 values for 8-bit coding, being only the first sixteenth of the QP range available to an 8-bit system.

5 It is thus useful to expand the range of QP in the general case. However, note that use of high precision intermediate encoding and/or decoding in conjunction with a direct correspondence between the range of QP and QP values (*i.e.*, a value x is the same as the representation number x ; thus, value 14 equals representation number 14) can result in high quality compression and decompression.

10 It is desirable to maintain the small number of steps in QP (such as 32 or 64 values, or some similar small number) if variation of QP within a frame is desired, since the bits required to code QP variations per macroblock are limited to 2 units in MPEG-4. If QP is varied only once per frame, or once per slice or similar large structure, then the number of values for QP can be large.

15 If needed, the effective range of QP can be extended to a wide range of determined values (*i.e.*, the values of QP actually applied during compression and decompression) within a small number of representational codes by taking advantage of the perceptually logarithmic nature of dynamic range. Accordingly, a non-linear lookup table can be used to map a small number of representational QP codes (such as the range of 1 to 32, or 1 to 128) into a much wider range of determined QP values (such as 1 to 128, 1 to 256, 1 to 1024, or 1 to 4096). In
20 such a table, the low QP code entries would map nearly one-to-one to determined QP values. For example, QP codes 1 to 4 might map to determined QP output values of 1 to 4. However, the mapping will gradually become steeper, in a logarithmic model, such that QP code 8 might map to determined QP value 16, QP code 16 might map to determined QP value 64, and QP code 32 might map to determined QP value 256. Note that MPEG-2 does provide for
25 a linear QP mode from 1-31, a double-step QP mode that maps each code from 1-31 to twice its value (*i.e.*, to 2-62), and a non-linear QP mode that maps codes 1-31 to determined values 1-112. In MPEG-2, these large determined QP values result in extremely crude coding using 8-bit pixel values. For example, QP values of 62 and 112 correspond to coding only two or one bits, respectively, out of the 8-bit values. Thus, any encoded image using these values
30 would be of extremely poor quality.

Alternatively, if QP is specified once per frame or slice or other large structure, the number of bits available for QP is not limited, and QP values can be fully represented over a very wide range within any appropriate number of bits, including 16 bits, 32 bits, 32 bit floating point, and even higher numbers of bits. However, the conceptual framework of wide
5 dynamic range images is such that some portions of the image are high in brightness, requiring high determined QP values, and other portions are low in brightness, requiring low determined QP values. Thus, it is useful to use a method for efficiently specifying QP values on a regional basis. The existing mechanism of coding systems (such as MPEG-4) of allowing QP values to vary ± 2 units per macroblock is sufficient if the range of QP codes is
10 limited (such as 1 to 32, as in MPEG-4). However, if a large QP value range is needed, other simple methods of specifying regional QP values are also appropriate and useful.

Thus, the use of regionally-varying QP values is sufficiently general to allow for very wide dynamic range representations which can be highly compressed, and yet be visually indistinguishable from the original moving image.

15 Attention must be paid to the number of bits available for the coded representation after dividing by determined QP values (dividing the DCT output by QP is also called "quantization"). After quantization, the remaining bits must be coded into the bitstream. The coded bits, except in the case of intra frames and intra macroblocks, represent the difference between the best motion-vector-predicted match in a previous or subsequent frame, and the
20 current frame and current macroblock. The coded representation of this quantized DCT-coded difference will determine the compression ratio that can be achieved.

In MPEG-2 and MPEG-4, the maximum coded value range is ± 2047 (limited by the VLC table representation). This corresponds to an unquantized precision of 8-bits. Thus, for unquantized (*i.e.*, QP=1) coding of 10-bit images, it is possible to exceed this maximum
25 coding range by a factor of four. This could happen if the best match predictor block contains a full-scale black to white transition in DC, or the equivalent full-scale AC coefficient change (such as a sharp black-white edge predicted from a flat-gray region). Optimal predictors rarely will provide so poor of a match, and thus full-scale coding will rarely, if ever, be required in this 10-bit example case. However, the range of coded values in a 12-bit moving
30 image, or in even wider 14 or 16-bit images, will often exceed a range limitation such as ± 2047 . While the ± 2047 limit is easily extended, it is conceptually beneficial to attempt to

limit the average number of bits being coded. Both the average number of bits, as well as the maximum coded range, are directly reduced by the use of QP values. For example, a QP of 4 extends the dynamic range available with a ± 2047 coded representation to include all possible cases of 10-bit coding, and all likely cases of 12-bit moving picture coding. It can therefore be seen that extending QP values to a higher number of values, such as 16, 32, or 64, can further extend the moving image dynamic range that can be represented by a limited number of coded bits, such as a ± 2047 range. Such higher QP values not only extend the range, but also reduce the average number of bits being coded, thus yielding high compression.

A key to this concept is that high QP values in wide dynamic range images correspond to image regions having high brightness, which do not require fine tonal distinction. This method of coding wide dynamic range images by maintaining extended precision for frames, and by utilizing high QP values in bright regions and low QP values in dark regions, can achieve high compression ratios with perceptually perfect coding which is indistinguishable from the original.

Thus, the utility of MPEG-like compression techniques can be extended for use in compressing moving wide-dynamic-range images. As cameras, projectors, and other image system components extend their range capabilities, this compression methodology can be applied to make highly efficient compression available. The resulting compression system is therefore extensible over many generations of future technology improvements in imaging systems.

In summary, one aspect of the invention includes application of high precision compression to wide dynamic range images by extending the range of QP values. The extension may be accomplished either by increasing the range of QP values directly, or indirectly through a non-linear transformation (such as a function or lookup table) which maps a small range of QP codes to a wider range of QP values for dividing (compression) and multiplying (decompression). Another aspect of the invention is the determination of such a wide-range QP values based on regional information, or by examination of information available during compression (*e.g.*, DCT coefficients, or the number of bits generated for a given number of candidate QP values, from which an appropriate one is selected), or a combination of the two determination methods.

Extended intermediate processing precision and an extended QP value range can also be applied to coding of both extended and reduced contrast range regions in a moving image. For example, it is common on hazy days to have high contrast on foreground objects, but have decreasing contrast with distance. Objects far away will often have very low contrast.

5 Other common situations, such as the scenes behind the windows of a building or a car windshield, also have reduced contrast due to the glass and reflection of the glass. The reflections also exhibit reduced contrast.

The principles of extended precision and QP value range can be applied to reduced contrast regions of an image to extend the precision with which the low-contrast portions are compression coded. As with dynamic range extension, which uses low QP values for dark

10 regions and high QP values for bright regions, low QP values can be utilized with low-contrast (often distant) regions, whereas high QP values are utilized for high contrast (usually foreground) regions. In this way, if a camera sensor has sufficient bit resolution, distant objects in the sky or on the ground on a hazy day can be fully distinguished when

15 decompressed. Their contrast can subsequently be artificially enhanced, thereby revealing a clear image having a normal dynamic range.

While current cameras and films are limited to approximately 10-bits of gamma or logarithmic dynamic range, future cameras quite likely will have higher precision. Such extended camera image brightness distinction would be useful for viewing detail in low

20 contrast areas, in addition to extending the dynamic range. As with cameras, as projectors extend their dynamic range and their maximum brightness, it is possible to distinguish low contrast details within this extended range. Subtle variations in brightness, such as falling raindrops, are much more easily seen on a wide dynamic range projector than on a limited dynamic range viewing monitor. An object displayed by a wide dynamic range projector is

25 easily distinguished because it has a wide range of brightness variation to the observer, whereas there is a low range of brightness variation on a computer CRT display. Thus, as cameras and displays expand their dynamic range and tonal distinction (*i.e.*, add additional bits of resolution, beyond the current 10-bit capabilities), it will be desirable to expand not only dynamic range but also contrast distinction.

30 Essentially the same techniques that support extended dynamic range also support high distinction coding of low contrast regions. In particular, QP values are applied to AC

coefficients of the DCT output in a manner distinct from the DC coefficient, which is usually coded specially (to improve DC term coding efficiency). The scaling of AC coefficients is therefore naturally ranged about the prevailing DC value. For example, a low contrast region in gray haze will have low amplitude AC coefficients about the DC haze average value.

5 Thus, applying low QP values will naturally preserve subtle tone variations within the low contrast regions. As with extended dynamic range, high QP values allow normal coding of high contrast foreground regions.

In order to adjust QP values to be appropriate for dark regions and hazy regions, yet still be suitable for normal contrast full brightness regions, QP values should be regionally
10 determined with respect to the contrast and brightness of each picture region. This can also be automatically determined if QP values are set for each macroblock such that each macroblock in a region generates approximately the same number of bits. For an image having wide dynamic range, including dark regions, as well as low contrast regions, giving each macroblock a constant number of bits will automatically optimize the representation
15 over the entire range of brightness and contrast. However, it is also desirable to provide more bits to high detail regions than to low detail regions, and to provide more bits to moving regions than static regions.

Determining a QP value for each macroblock can be automated by examining the relative amplitudes of the DCT coefficients in each macroblock. Macroblocks containing
20 DCT coefficients that indicate detail and those that indicate motion can be provided more bits than those macroblocks where the relative weights of coefficients indicate either low detail or low change (motion). However, the noise of the camera sensor must also be taken into account, since noise will behave like both change (motion) and detail (high frequency coefficients). When used with a true wide dynamic range and high distinction sensor of
25 suitable low noise level, the DCT coefficient relative weightings themselves can form an appropriate indicator for automatically setting the value of QP. In particular, larger DCT coefficients yield larger QP values. Accordingly, a mapping or correlation between DCT coefficients and desired corresponding QP values can be empirically determined.

Simple regional algorithms, such as maximum region brightness and contrast, are
30 another method that can be used to determine (or help determine, together with other mechanisms) appropriate QP values. Additional information can also be provided by regional

detail amplitude (picture high frequency static) measurement algorithms. Each method has its own advantages. However, the DCT coefficients are themselves sufficient to determine QP values in the case of intra macroblocks. That is, the DCT coefficients are a measure of detail combined with motion for predicted macroblocks, so the use of a separate detail measure
5 (such as a parallel intra DCT transform) can help isolate motion detail changes (such as rain drops or moving water waves on the horizon) from the detail of the current frame image macroblock (static after motion compensation, such as blades of grass with a slowly moving camera). The simplicity of the use of the DCT coefficients themselves to indicate QP makes it a particularly attractive method for practical implementation.

10 FIG. 5 is a flowchart that summarizes a preferred method of extending dynamic range and/or contrast during image compression.

Step 500: Begin with a wide dynamic range or wide contrast range source picture.

Step 502: If needed, extend the effective value range of the quantization parameter (QP) code set. This may be done, for example, by one of the following techniques:

15 1) Extend the set of QP codes from a nominal range (typically 32 levels) to a larger actual range (128, 1024, or 4096 levels, or whatever is appropriate for the image range). Thus, the nominal values directly represent an extended range of possible values.

2) Use a non-linear lookup table or mapping function to correlate nominal QP codes in a non-linear way to a larger effective range of values. The mapping typically would
20 be linear at low values, but increase in effective QP multiply and divide step size as values increase toward a typical range maximum. For example, 32 or 64 codes may be expanded using a non-linear lookup or mapping function to yield a larger effective range having a larger maximum value, such as 128, 1024, 4096, or whatever is appropriate for the image range.

25 *Step 504:* Determine the QP value that should be coded for each macroblock of an image undergoing compression, preferably using one of the following methods:

1) Determine an appropriate QP value using algorithms to determine local image region contrast within each frame, local dynamic range within each frame, local detail amplitudes within each frame, and local motion between one or more sequential frames (as
30 described above), from an analysis of the moving image stream

2) Determine a QP value based upon information generated in the compression process, based upon number of bits generated (for a number of candidate values of QP), and the amplitude and frequency of DCT coefficients prior to quantization.

3) Apply a combination of the information from 1) and 2), determining a QP value for each macroblock utilizing both regional information as well as information generated from the compression process.

Step 506: Use extended precision for all intermediate processing, as described above, to compress the image using the determined QP value(s) from Step 504. The compressed image, along with the associated nominal QP codes corresponding to the determined QP values used during compression, may be stored or transmitted, as desired.

Step 508: Decompress the stored or transmitted image, using high precision decompression as described above, to a wide-dynamic range, wide contrast range, high resolution image for various applications. The associated nominal QP codes are mapped, if necessary, back to corresponding determined QP values for such decompression. Such applications include home and theatrical presentation of movies and sports, archiving of stored images, business uses of moving image presentations, government applications (*e.g.*, surveillance, military command and control), *etc.* The decompressed images can be viewed on wide-dynamic range display devices and/or used as a source for image analysis using algorithms which benefit from (or require) high quality wide-dynamic range images in order to provide optimal analysis (such algorithms are not the subject of this disclosure).

Implementation

The invention may be implemented in hardware or software, or a combination of both (*e.g.*, programmable logic arrays). Unless otherwise specified, the algorithms included as part of the invention are not inherently related to any particular computer or other apparatus. In particular, various general purpose machines may be used with programs written in accordance with the teachings herein, or it may be more convenient to construct more specialized apparatus (*e.g.*, integrated circuits) to perform the required method steps. Thus, the invention may be implemented in one or more computer programs executing on one or more programmable computer systems each comprising at least one processor, at least one data storage system (including volatile and non-volatile memory and/or storage elements), at

least one input device or port, and at least one output device or port. Program code is applied to input data to perform the functions described herein and generate output information. The output information is applied to one or more output devices, in known fashion.

5 Each such program may be implemented in any desired computer language (including machine, assembly, or high level procedural, logical, or object oriented programming languages) to communicate with a computer system. In any case, the language may be a compiled or interpreted language.

10 Each such computer program is preferably stored on or downloaded to a storage media or device (*e.g.*, solid state memory or media, or magnetic or optical media) readable by a general or special purpose programmable computer, for configuring and operating the computer when the storage media or device is read by the computer system to perform the procedures described herein. The inventive system may also be considered to be implemented as a computer-readable storage medium, configured with a computer program, where the storage medium so configured causes a computer system to operate in a specific and predefined manner to perform the functions described herein.

15 A number of embodiments of the invention have been described. Nevertheless, it will be understood that various modifications may be made without departing from the spirit and scope of the invention. For example, some of the steps described above may be order independent, and thus can be performed in an order different from that described. Accordingly, other embodiments are within the scope of the following claims.

WHAT IS CLAIMED IS:

1 1. A method for compressing a sequence of digitized video images including a sequence of
2 frames represented at a first precision in a first color space, the method including:
3 (a) transforming the sequence of frames to a representation in a second color space at a
4 second precision greater than the first precision; and
5 (b) performing subsequent encoding steps at the second precision to create a
6 compressed output.

1 2. A method for compressing a sequence of digitized video images including a sequence of
2 frames represented at a first precision in a first color space, the method including:
3 (a) transforming the sequence of frames to a representation in a second color space at a
4 second precision greater than the first precision;
5 (b) performing a motion compensated discrete cosine transform at the second precision
6 on the sequence of frames to produce a first encoded output;
7 (c) quantizing the first encoded output at the second precision to create a quantized
8 output; and
9 (d) performing an inverse discrete cosine transform at the second precision on the
10 quantized output to produce a compressed output at the second precision.

1 3. A method for compressing and decompressing a sequence of digitized video images
2 including a sequence of frames represented at a first precision, the method including
3 encoding the sequence of frames at a second precision greater than the first precision to
4 create a compressed bitstream.

1 4. The method of claim 3, further including decoding the compressed bitstream at the
2 second precision.

1 5. The method of claim 3, wherein encoding the sequence of frames includes transforming
2 the sequence of frames to a representation in a second color space at the second precision.

1

- 1 6. A method for compressing and decompressing a sequence of digitized video images
2 including a sequence of frames represented at a first precision in a first color space, the
3 method including:
 - 4 (a) transforming the sequence of frames to a representation in a second color space at a
5 second precision greater than the first precision;
 - 6 (b) performing subsequent encoding steps at the second precision to create a
7 compressed bitstream; and
 - 8 (c) decoding the compressed bitstream by:
 - 9 (1) dequantizing the compressed bitstream at the second precision to create a
10 dequantized output;
 - 11 (2) applying an inverse discrete cosine transform at the second precision on the
12 dequantized output to produce a decompressed output;
 - 13 (3) generating image frames at the second precision from the decompressed output.
- 1 7. The method of claim 6, wherein the subsequent encoding steps include applying an
2 inverse discrete cosine transform numerical algorithm, and wherein decoding the
3 compressed bitstream includes applying a matching inverse discrete cosine transform
4 numerical algorithm on the dequantized output.
- 1 8. The method of claim 6, further including generating P frames at the second precision.
- 1 9. The method of claim 6, further including generating B frames at the second precision.

1 10. A method for compressing and decompressing a sequence of digitized video images
2 including a sequence of frames represented at a first precision in a first color space, the
3 method including:

- 4 (a) transforming the sequence of frames to a representation in a second color space at a
5 second precision greater than the first precision;
6 (b) performing subsequent encoding steps at the second precision to create a
7 compressed bitstream, including applying an inverse discrete cosine transform
8 numerical algorithm; and
9 (c) performing subsequent decoding steps at the second precision on the compressed
10 bitstream, including applying a matching inverse discrete cosine transform
11 numerical algorithm.

1 11. A method for increasing the average compression ratio during compression of a sequence
2 of digitized video images, the digitized video images including a sequence of frames, into
3 a compressed bitstream that includes I frames, the method including encoding the
4 sequence of frames by applying an inverse discrete cosine transform numerical algorithm
5 that matches the inverse discrete cosine transform numerical algorithm applied during a
6 subsequent decoding process, thereby generating a compressed bitstream having
7 increased spacing between I frames.

1 12. The method of claim 11, further including decoding the compressed bitstream, including
2 applying an inverse discrete cosine transform numerical algorithm that matches the
3 inverse discrete cosine transform numerical algorithm applied during encoding.

1

- 1 13. A method for compressing a sequence of digitized video images having at least one of a
2 wide dynamic range or wide contrast range, the video images including frames
3 represented at a first precision, each frame including a plurality of macroblocks defining
4 regions, the method including:
- 5 (a) determining a quantization parameter, QP, having codes that represent either a direct
6 correspondence to possible values or an extended range of possible values;
 - 7 (b) determining a QP value for each macroblock of each frame;
 - 8 (c) compressing each frame at a second precision higher than the first precision to
9 create a compressed frame, such compressing including applying the determined QP
10 values for such frame to reduce the number of bits required to encode such frame;
 - 11 (d) associating QP codes with the determined QP values used during compressing; and
12 (e) outputting each compressed frame and the associated QP codes.
- 1 14. The method of claim 13, further including directly extending the range of possible QP
2 values by increasing the number of bits representing QP codes.
- 1 15. The method of claim 13, further including effectively extending the range of possible QP
2 values by mapping QP codes to a wider range of possible QP values.
- 1 16. The method of claim 13, wherein determining a QP value for each macroblock of each
2 frame is based upon regional information with respect to one or more frames.
- 1 17. The method of claim 16, wherein the regional information includes local image region
2 contrast within each frame.
- 1 18. The method of claim 16, wherein the regional information includes local dynamic range
2 within each frame.
- 1 19. The method of claim 16, wherein the regional information includes local detail
2 amplitudes within each frame.

1 20. The method of claim 16, wherein the regional information includes local motion between
2 one or more sequential frames.

1 21. The method of claim 13, wherein determining a QP value for each macroblock of each
2 frame is based upon information generated during the compressing step.

1 22. The method of claim 21, wherein the information generated during the compressing step
2 includes relative amplitudes of discrete cosine transform coefficients for each
3 macroblock.

1 23. The method of claim 21, wherein the information generated during the compressing step
2 includes a constant number of bits assigned to encode each macroblock within a region of
3 a frame.

1 24. The method of claim 13, further including:

2 (a) for each compressed frame, re-determining a QP value for each QP code associated
3 with such compressed frame; and

4 (b) decompressing each compressed frame at the second precision to create a
5 decompressed frame having at least one of a wide dynamic range or wide contrast
6 range, such decompressing including applying the re-determined QP values for such
7 frame.

1

1 25. A computer program, stored on a computer-readable medium, for compressing a
2 sequence of digitized video images including a sequence of frames represented at a first
3 precision in a first color space, the computer program comprising instructions for causing
4 a computer to:

- 5 (a) transform the sequence of frames to a representation in a second color space at a
6 second precision greater than the first precision; and
7 (b) perform subsequent encoding steps at the second precision to create a compressed
8 output.

1 26. A computer program, stored on a computer-readable medium, for compressing a
2 sequence of digitized video images including a sequence of frames represented at a first
3 precision in a first color space, the computer program comprising instructions for causing
4 a computer to:

- 5 (a) transform the sequence of frames to a representation in a second color space at a
6 second precision greater than the first precision;
7 (b) perform a motion compensated discrete cosine transform at the second precision on
8 the sequence of frames to produce a first encoded output;
9 (c) quantize the first encoded output at the second precision to create a quantized
10 output; and
11 (d) perform an inverse discrete cosine transform at the second precision on the
12 quantized output to produce a compressed output at the second precision.

1

1 27. A computer program, stored on a computer-readable medium, for compressing and
2 decompressing a sequence of digitized video images including a sequence of frames
3 represented at a first precision in a first color space, the computer program comprising
4 instructions for causing a computer to encode the sequence of frames at a second
5 precision greater than the first precision to create a compressed bitstream.

1 28. The computer program of claim 27, further including instructions for causing the
2 computer to decode the compressed bitstream at the second precision.

1 29. The computer program of claim 27, wherein the instructions for causing the computer to
2 encode the sequence of frames includes instructions for causing the computer to
3 transform the sequence of frames to a representation in a second color space at the second
4 precision.

1

1 30. A computer program, stored on a computer-readable medium, for compressing and
2 decompressing a sequence of digitized video images including a sequence of frames
3 represented at a first precision in a first color space, the computer program comprising
4 instructions for causing a computer to:
5 (a) transform the sequence of frames to a representation in a second color space at a
6 second precision greater than the first precision;
7 (b) perform subsequent encoding steps at the second precision to create a compressed
8 bitstream; and
9 (c) decode the compressed bitstream by:
10 (1) dequantizing the compressed bitstream at the second precision to create a
11 dequantized output;
12 (2) applying an inverse discrete cosine transform at the second precision on the
13 dequantized output to produce a decompressed output;
14 (3) generating image frames at the second precision from the decompressed output.

1 31. The computer program of claim 30, wherein the subsequent encoding steps include
2 instructions for causing the computer to apply an inverse discrete cosine transform
3 numerical algorithm, and wherein the instructions for causing the computer to decode the
4 compressed bitstream include instructions for causing the computer to apply a matching
5 inverse discrete cosine transform numerical algorithm on the dequantized output.
1

1 32. The computer program of claim 31, further including instructions for causing the
2 computer to generate P frames at the second precision.

1 33. The computer program of claim 31, further including instructions for causing the
2 computer to generate B frames at the second precision.

1 34. A computer program, stored on a computer-readable medium, for compressing and
2 decompressing a sequence of digitized video images including a sequence of frames
3 represented at a first precision in a first color space, the computer program comprising
4 instructions for causing a computer to:

- 5 (a) transform the sequence of frames to a representation in a second color space at a
6 second precision greater than the first precision;
- 7 (b) perform subsequent encoding steps at the second precision to create a compressed
8 bitstream, including applying an inverse discrete cosine transform numerical
9 algorithm; and
- 10 (c) perform subsequent decoding steps at the second precision on the compressed
11 bitstream, including applying a matching inverse discrete cosine transform
12 numerical algorithm.

1 35. A computer program, stored on a computer-readable medium, for increasing the average
2 compression ratio during compression of a sequence of digitized video images, the
3 digitized video images including a sequence of frames, into a compressed bitstream that
4 includes I frames, the computer program comprising instructions for causing a computer
5 to encode the sequence of frames by applying an inverse discrete cosine transform
6 numerical algorithm that matches the inverse discrete cosine transform numerical
7 algorithm applied during a subsequent decoding process, thereby generating a
8 compressed bitstream having increased spacing between I frames.

1 36. The computer program of claim 35, further including instructions for causing the
2 computer to decode the compressed bitstream, including instructions for causing the
3 computer to apply an inverse discrete cosine transform numerical algorithm that matches
4 the inverse discrete cosine transform numerical algorithm applied during encoding.

1 37. A computer program, stored on a computer-readable medium, for compressing a
2 sequence of digitized video images having at least one of a wide dynamic range or wide
3 contrast range, the video images including frames represented at a first precision, each
4 frame including a plurality of macroblocks defining regions, the computer program
5 comprising instructions for causing a computer to:

- 6 (a) determine a quantization parameter, QP, having codes that represent either a direct
7 correspondence to possible values or an extended range of possible values;
- 8 (b) determine a QP value for each macroblock of each frame;
- 9 (c) compress each frame at a second precision higher than the first precision to create a
10 compressed frame, such compressing including applying the determined QP values
11 for such frame to reduce the number of bits required to encode such frame;
- 12 (d) associate QP codes with the determined QP values used during compressing; and
13 (e) output each compressed frame and the associated QP codes.

1 38. The computer program of claim 37, further including instructions for causing the
2 computer to directly extend the range of possible QP values by increasing the number of
3 bits representing QP codes.

1 39. The computer program of claim 37, further including instructions for causing the
2 computer to effectively extend the range of possible QP values by mapping QP codes to a
3 wider range of possible QP values.

1 40. The computer program of claim 37, wherein determining a QP value for each macroblock
2 of each frame is based upon regional information with respect to one or more frames.

1 41. The computer program of claim 40, wherein the regional information includes local
2 image region contrast within each frame.

- 1 42. The computer program of claim 40, wherein the regional information includes local
2 dynamic range within each frame.
- 1 43. The computer program of claim 40, wherein the regional information includes local detail
2 amplitudes within each frame.
- 1 44. The computer program of claim 40, wherein the regional information includes local
2 motion between one or more sequential frames.
- 1 45. The computer program of claim 37, wherein determining a QP value for each macroblock
2 of each frame is based upon information generated during the compressing step.
- 1 46. The computer program of claim 45, wherein the information generated during the
2 compressing step includes relative amplitudes of discrete cosine transform coefficients
3 for each macroblock.
- 1 47. The computer program of claim 45, wherein the information generated during the
2 compressing step includes a constant number of bits assigned to encode each macroblock
3 within a region of a frame.
- 1 48. The computer program of claim 47, further including instructions for causing the
2 computer to:
3 (a) for each compressed frame, re-determine a QP value for each QP code associated
4 with such compressed frame; and
5 (b) decompress each compressed frame at the second precision to create a
6 decompressed frame having at least one of a wide dynamic range or wide contrast
7 range, such decompression including applying the re-determined QP values for such
8 frame.

1 49. A system for compressing a sequence of digitized video images including a sequence of
2 frames represented at a first precision in a first color space, the system including:
3 (a) means for transforming the sequence of frames to a representation in a second color
4 space at a second precision greater than the first precision; and
5 (b) means for performing subsequent encoding steps at the second precision to create a
6 compressed output.

1 50. A system for compressing a sequence of digitized video images including a sequence of
2 frames represented at a first precision in a first color space, the system including:
3 (a) means for transforming the sequence of frames to a representation in a second color
4 space at a second precision greater than the first precision;
5 (b) means for performing a motion compensated discrete cosine transform at the second
6 precision on the sequence of frames to produce a first encoded output;
7 (c) means for quantizing the first encoded output at the second precision to create a
8 quantized output; and
9 (d) means for performing an inverse discrete cosine transform at the second precision
10 on the quantized output to produce a compressed output at the second precision.

1 51. A system for compressing and decompressing a sequence of digitized video images
2 including a sequence of frames represented at a first precision in a first color space, the
3 system including means for encoding the sequence of frames at a second precision
4 greater than the first precision to create a compressed bitstream.

1 52. The system of claim 51, further including means for decoding the compressed bitstream
2 at the second precision.

1 53. The system of claim 51, wherein the means for encoding the sequence of frames includes
2 means for transforming the sequence of frames to a representation in a second color space
3 at the second precision.

1

1 54. A system for compressing and decompressing a sequence of digitized video images
2 including a sequence of frames represented at a first precision in a first color space, the
3 system including:
4 (a) means for transforming the sequence of frames to a representation in a second color
5 space at a second precision greater than the first precision;
6 (b) means for performing subsequent encoding steps at the second precision to create a
7 compressed bitstream; and
8 (c) means for decoding the compressed bitstream by:
9 (1) dequantizing the compressed bitstream at the second precision to create a
10 dequantized output;
11 (2) applying an inverse discrete cosine transform at the second precision on the
12 dequantized output to produce a decompressed output;
13 (3) generating image frames at the second precision from the decompressed output.

1 55. The system of claim 54, wherein the means for performing subsequent encoding steps
2 include means for applying an inverse discrete cosine transform numerical algorithm, and
3 wherein the means for decoding the compressed bitstream includes means for applying a
4 matching inverse discrete cosine transform numerical algorithm on the dequantized
5 output.

1 56. The system of claim 54, further including means for generating P frames at the second
2 precision.

1 57. The system of claim 54, further including means for generating B frames at the second
2 precision.

1 58. A system for compressing and decompressing a sequence of digitized video images
2 including a sequence of frames represented at a first precision in a first color space, the
3 system including:
4 (a) means for transforming the sequence of frames to a representation in a second color
5 space at a second precision greater than the first precision;
6 (b) means for performing subsequent encoding steps at the second precision to create a
7 compressed bitstream, including applying an inverse discrete cosine transform
8 numerical algorithm; and
9 (c) means for performing subsequent decoding steps at the second precision on the
10 compressed bitstream, including applying a matching inverse discrete cosine
11 transform numerical algorithm.

1 59. A system for increasing the average compression ratio during compression of a sequence
2 of digitized video images, the digitized video images including a sequence of frames, into
3 a compressed bitstream that includes I frames, the system including means for encoding
4 the sequence of frames by applying an inverse discrete cosine transform numerical
5 algorithm that matches the inverse discrete cosine transform numerical algorithm applied
6 during a subsequent decoding process, thereby generating a compressed bitstream having
7 increased spacing between I frames.

1 60. The system of claim 59, further including means for decoding the compressed bitstream,
2 including means for applying an inverse discrete cosine transform numerical algorithm
3 that matches the inverse discrete cosine transform numerical algorithm applied during
4 encoding.
1

1 61. A system for compressing a sequence of digitized video images having at least one of a
2 wide dynamic range or wide contrast range, the video images including frames
3 represented at a first precision, each frame including a plurality of macroblocks defining
4 regions, the system including:

5 (a) means for determining a quantization parameter, QP, having codes that represent
6 either a direct correspondence to possible values or an extended range of possible
7 values;

8 (b) means for determining a QP value for each macroblock of each frame;

9 (c) means for compressing each frame at a second precision higher than the first
10 precision to create a compressed frame, including means for applying the
11 determined QP values for such frame to reduce the number of bits required to
12 encode such frame;

13 (d) means for associating QP codes with the determined QP values used during
14 compressing; and

15 (e) means for outputting each compressed frame and the associated QP codes.

1 62. The system of claim 61, further including means for directly extending the range of
2 possible QP values by increasing the number of bits representing QP codes.

1 63. The system of claim 61, further including means for effectively extending the range of
2 possible QP values by mapping QP codes to a wider range of possible QP values.

1 64. The system of claim 61, wherein the means for determining a QP value for each
2 macroblock of each frame includes means for making such a determination based upon
3 regional information with respect to one or more frames.

1 65. The system of claim 64, wherein the regional information includes local image region
2 contrast within each frame.

1 66. The system of claim 64, wherein the regional information includes local dynamic range
2 within each frame.

1 67. The system of claim 64, wherein the regional information includes local detail amplitudes
2 within each frame.

1 68. The system of claim 64, wherein the regional information includes local motion between
2 one or more sequential frames.

1 69. The system of claim 61, wherein the means for determining a QP value for each
2 macroblock of each frame includes means for making such a determination based upon
3 information generated during the compressing step.

1 70. The system of claim 69, wherein the information generated by the means for compressing
2 includes relative amplitudes of discrete cosine transform coefficients for each
3 macroblock.

1 71. The system of claim 69, wherein the information generated by the means for compressing
2 includes a constant number of bits assigned to encode each macroblock within a region of
3 a frame.

1 72. The system of claim 61, further including:

2 (a) means for re-determining a QP value for each QP code associated with each
3 compressed frame; and

4 (b) means for decompressing each compressed frame at the second precision to create a
5 decompressed frame having at least one of a wide dynamic range or wide contrast
6 range, such decompressing including applying the re-determined QP values for such
7 frame.

1

ABSTRACT

Methods, systems, and computer programs for improved quality video compression.

Image quality from MPEG-style video coding may be improved by preserving a higher number of bits during intermediate encoding and decoding processing steps. Problems of inverse discrete cosine transform (IDCT) mismatch can be eliminated by exactly matching the IDCT function numerical algorithm of the decoder to the IDCT function numerical algorithm used for the decoding portion of the encoder. Also included is an application of high precision compression to wide dynamic range images by extending the range of the "quantization parameter" or "QP".

The extension of QP may be accomplished either by increasing the range of QP directly, or indirectly through a non-linear transformation. Also included is an application of extended intermediate processing precision and an extended QP range to reduced contrast regions of an image to extend the precision with which the low-contrast portions are compression coded.

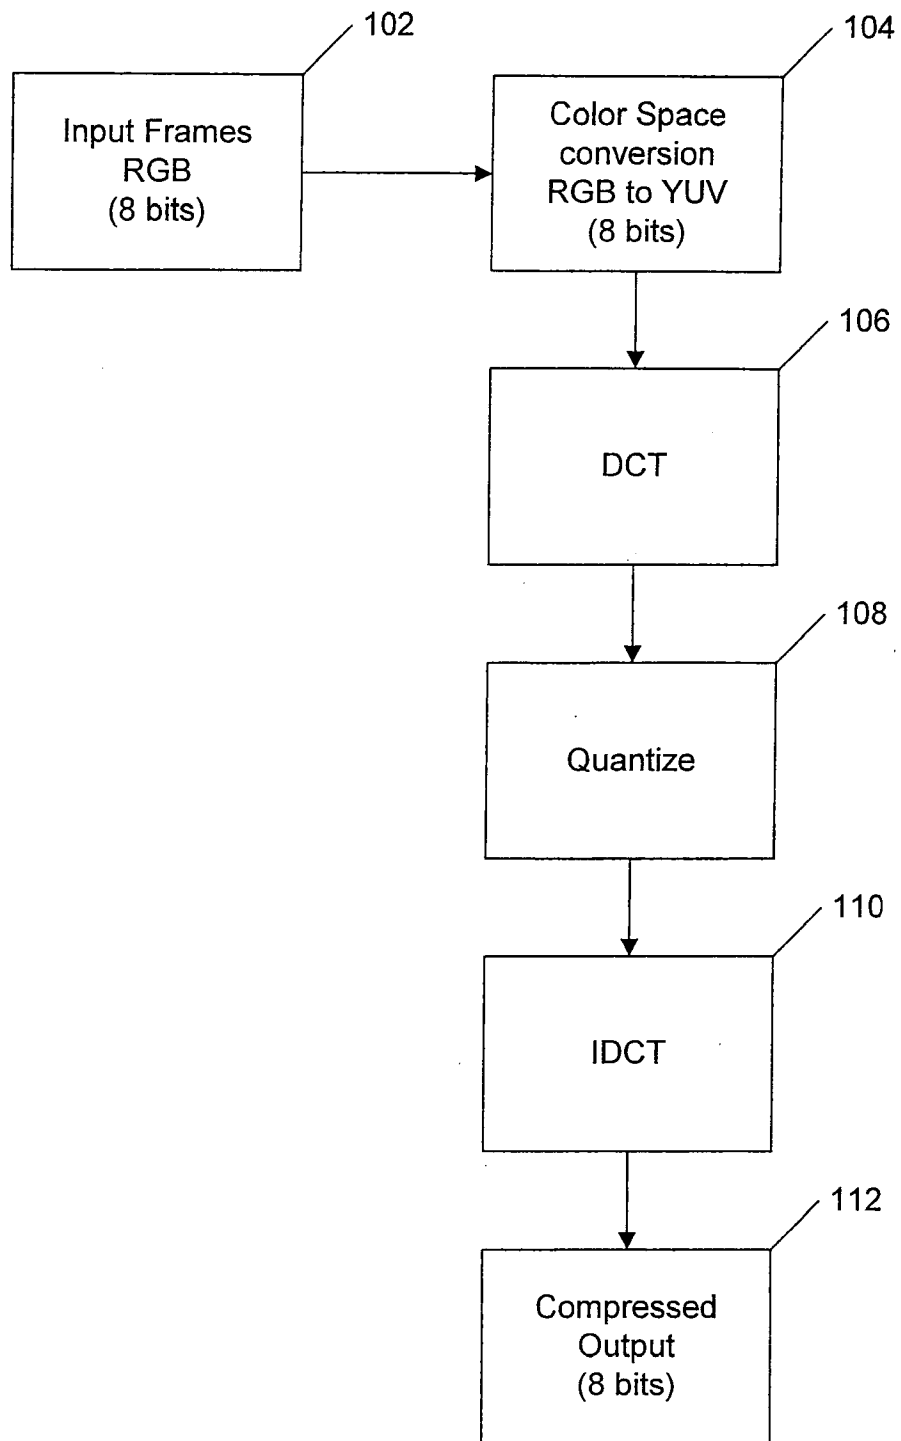


FIG. 1
Prior art MPEG-2 encoder

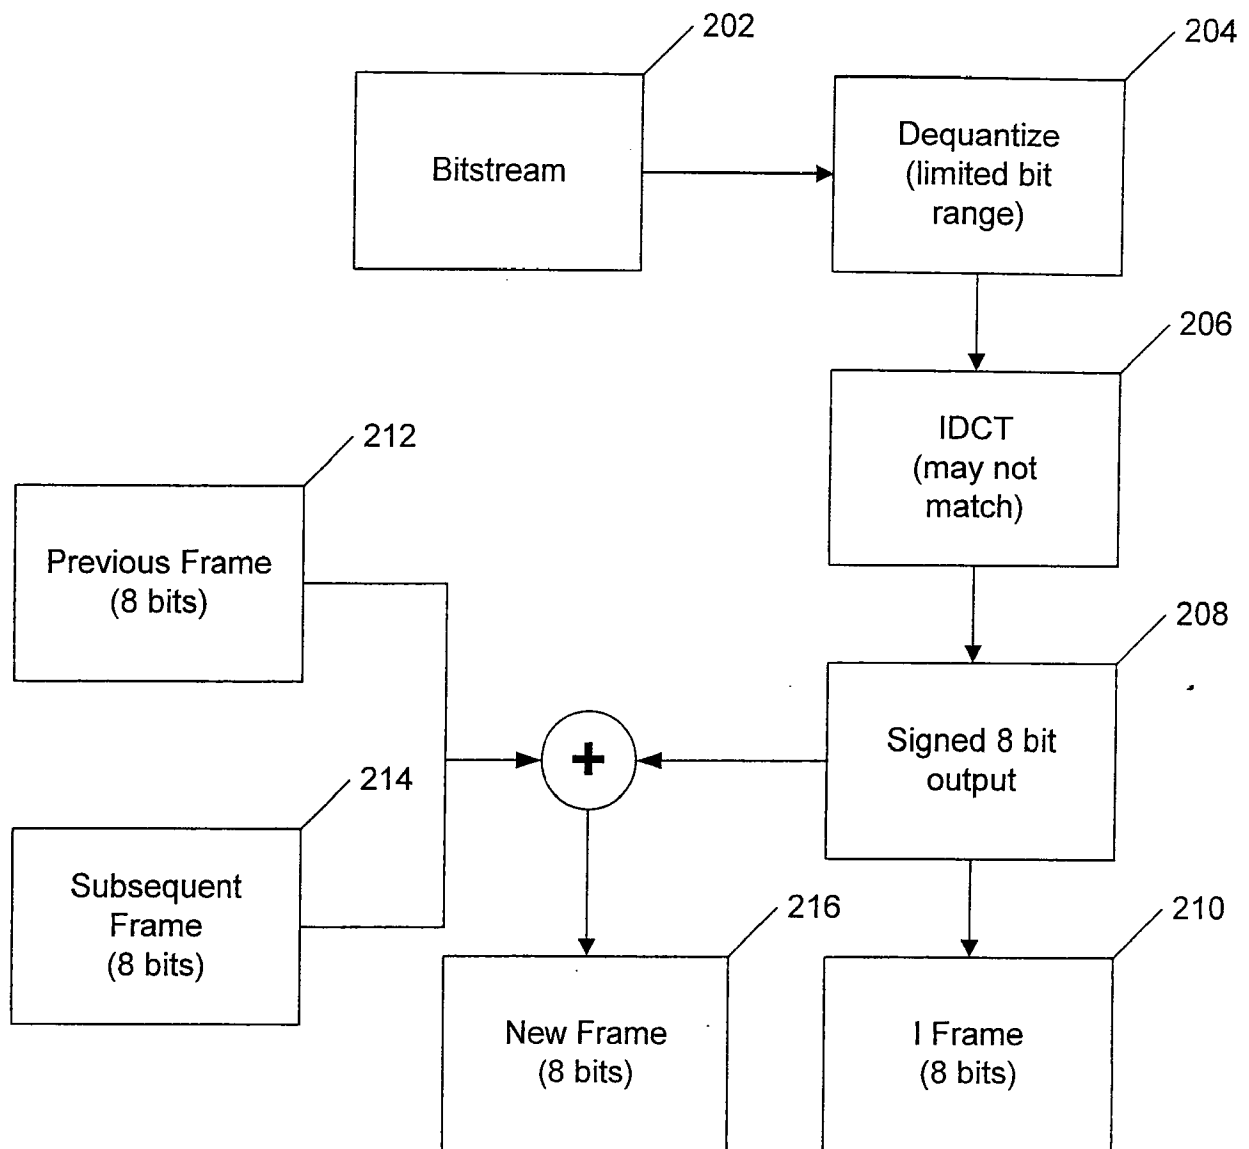


FIG. 2
Prior art MPEG-2 decoder




FIG. 3

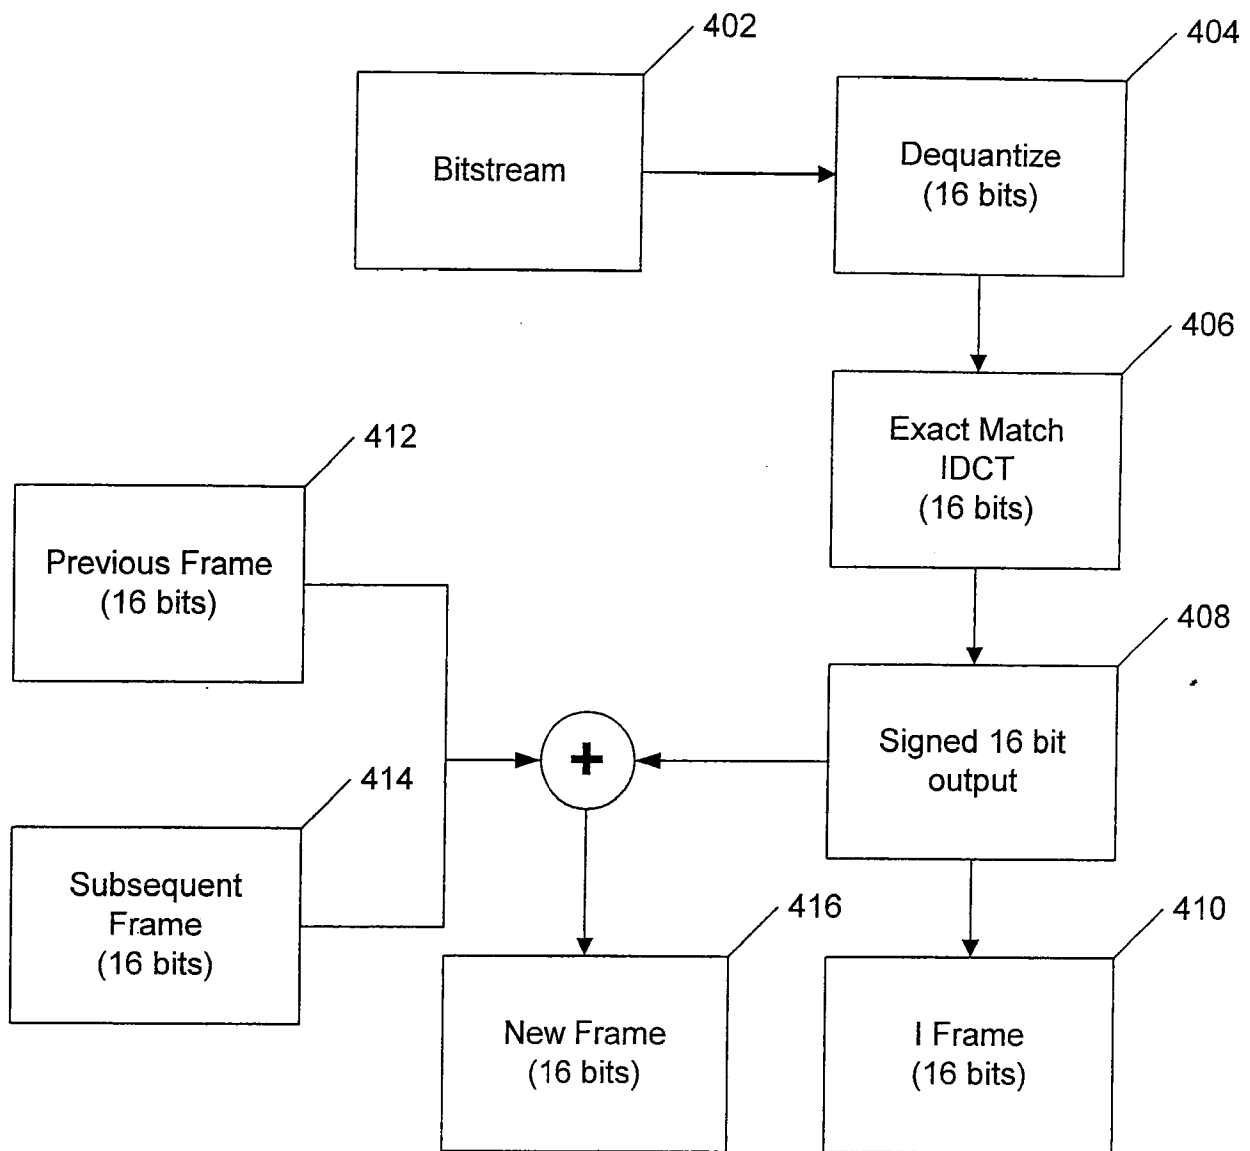


FIG. 4

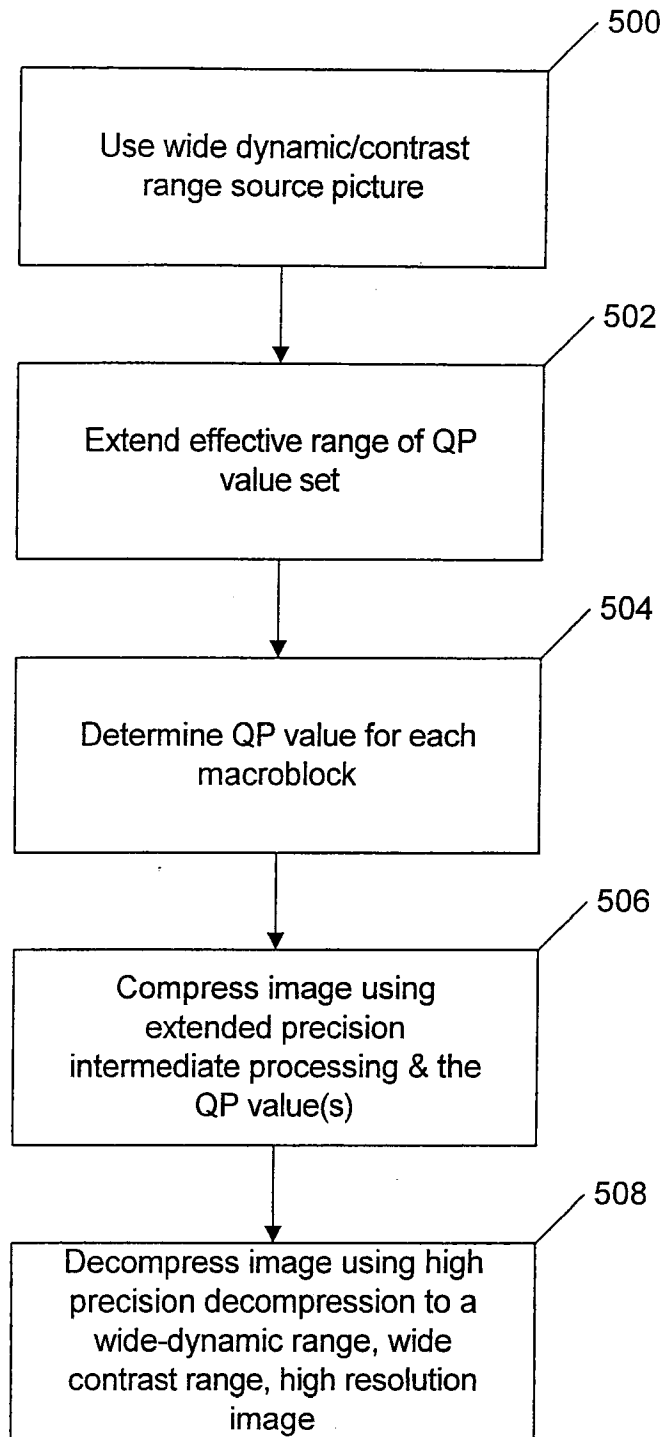


FIG. 5